

Relating Risk and Reliability Predictions to Design and Development Choices

Martin S. Feather, Ph. D., Jet Propulsion Laboratory, California Institute of Technology
Steven L. Cornford, Ph. D., Jet Propulsion Laboratory, California Institute of Technology

Key Words: Risk, Reliability, Design, Optimization

SUMMARY & CONCLUSIONS

Ideally, planning and managing the design and development of a complex systems should:

- Consider the *entire* lifecycle (design, development, testing, integration, deployment, operation and decommissioning).
- Take *risk* and *reliability* into account, as well as more traditional measures such as cost and performance, when making tradeoff decisions.
- Provide understanding, and therefore motivation, as to the *purpose* of design artifacts (that piece is there to...) and development activities (we'll be performing this test to ...).

A key enabler to all of the above is the ability to relate design and development *choices* to risk and reliability *predictions*. If the number of choices were small (e.g., selection between a handful of alternatives) it would suffice to perform reliability analysis on them individually, and view the results side-by-side. The challenge is that in many cases the number of design and development alternatives is large. In this context, the problem of relating reliability predictions to design and development choices is non-trivial.

Ongoing work towards a solution to this problem is the focus of this paper. Over several years we have developed a risk-based model the hallmark of which is the explicit representation of risk mitigations as options. We describe how this model functions, and the major implications of making mitigation options first class objects within an (otherwise relatively simple) analysis model. We also describe elaborations to this model's representation of risks, notably by the incorporation of fault tree notions. These improve the fidelity of the designs we are able to represent, and also offer the ability to represent design alternatives within the same framework. Finally, we describe the connections we are building between our risk analysis tool and other risk tools. The latter have greater strengths in their ability to represent and calculate over more elaborate risk structures, while our approach lends to them the aforementioned explicit treatment of the various forms of options for risk reduction.

1. INTRODUCTION

Our work originated as a method intended for planning the quality assurance of hardware systems [1]. In this context

there are many possible assurance activities. Some focus on the *prevention* of defects – for example, up-front planning, adoption of design standards, configuration management, training, etc. Others focus on the *detection* of defects – either to detect latent defects in a system (and so be able to correct them before actual deployment of the system), or to increase confidence that such defects are not present. For example, a wide gamut of reviews, design walkthroughs, tests, inspections, analyses, etc. can be applied to systems and their components.

Generally the total costs (e.g., time and budget) were *all* these activities and practices to be adopted would far exceed the resources available. In order to help in planning *which* of these activities and practices to adopt, Cornford developed the “Defect Detection and Prevention” (DDP) process. It treats assurance activities and practices as *options*, each of which are linked to the kinds of defects they prevent or detect. Each option-defect link is accompanied by a *quantitative* measure, of *effect* – the proportion by which the option will prevent or detect reduce the defect. These defects are in turn linked to the system-level objectives that they threaten. Again, each defect-objective link is accompanied by a *quantitative* measure, of *impact* – the proportion by which the defect, were it present, would detract from attainment of the objective. This scheme is illustrated abstractly in Figure 1.

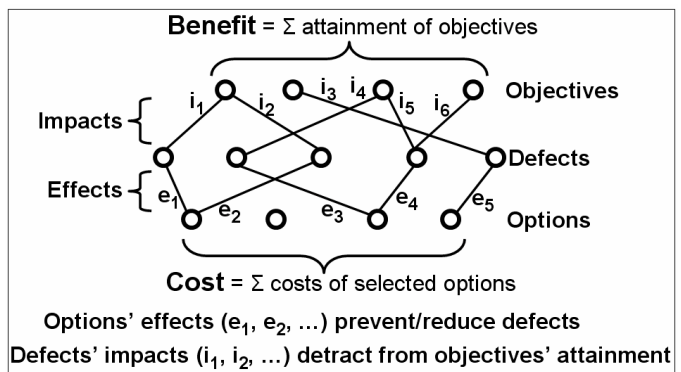


Figure 1 – Influence diagram of DDP concepts

The net result is the ability to estimate both the *cost* of a selection of options (what resources it would take to perform that selection), and the *benefit* of that selection (in terms of attainment of objectives). The latter is calculated by taking into account the defect reducing effects (preventions and detections) of the selected options, and in turn the (reduced)

extent to which the defects that evaded prevention or detection would have on attainment of objectives.

Cornford's initial experiments in 1998 showed the merits of this approach in the area of assurance planning, where options are choices among the gamut of assurance-related activities. It makes apparent situations where the selection of options is less than ideal – for example, where overly-many options are selected to reduce the same defect, while other defects go relatively unaddressed. It also clarifies the purpose of the selected options – namely as preventions or detections of the kinds of defects to which they are linked. These two areas of insight help improve both the planning and conduct of assurance activities, by helping to arrive at a cost-effective selection of them, and by clarifying the purpose(s) of the individual assurance activities themselves. In these experiments spreadsheets were used to store and calculate with the quantitative information. These proved sufficient to demonstrate the value of the approach, but suggested ways in which custom software should be developed to enable the process to proceed more smoothly. Such software has since been developed [2]. In the years since we have applied this to a wide range of problems, most often to systems and technologies at a relatively early stage in their development lifecycles to help plan their subsequent development.

The relevance of this approach to risk and reliability stems from its probabilistic treatment of defects, using the traditional decision statistic risk factor of expected loss [3]. It treats a defect as having a likelihood of occurrence (which may be decreased by adoption of the assurance options), and impacts on various objectives. The sum total of those impacts is equivalent to the traditional risk notion of “severity” (a.k.a. “consequence”). In subsequent work, we have broadened the area of application to early-phase project planning, and for such purposes often switch terminology accordingly, using “failure mode” or “risk” in place of “defect”, and “mitigation” in place of “option”. Indeed, this approach is very much motivated by techniques drawn from the Probabilistic Risk Assessment (PRA) community (e.g., for an overview, see [4]). It is the explicit treatment of (defect- or risk-reducing) options that sets this work apart from traditional PRA.

In the sections that follow we investigate issues that arise from making explicit the options for defect-reduction (or equivalently, risk-reduction), and the work we are performing to extend the model to incorporate (or connect to other tools with) more elaborate representations of risk.

2. IMPLICATIONS OF MAKING OPTIONS EXPLICIT

The focus of most of the mainstream work on risk and reliability has been on methods to assess the risk or reliability of a given design. Adding into these approaches the *explicit* treatment of defect- or risk-reducing options has benefits; these are outlines, along with what it takes to realize them.

2.1 Benefits of making options explicit

The main benefits of adding into these approaches the *explicit* treatment of defect- or risk-reducing options are as follows:

Ability to incorporate consideration of design practices when performing risk and reliability calculations. For example, the effects of preventative measures can be assessed in terms of their reduction on the prevalence of defects; test, analyses, inspections, reviews, etc., can be assessed in terms of their ability to detect defects (in advance of deployment and operation of the system, and therefore in time to correct them).

Ability to perform tradeoff decisions that take risk into account. In the DDP approach, risk is an intermediate concept – intermediate between the objectives whose attainments it threatens, and the options whose adoption (at a cost) decreases risk. The overall DDP model allows for the computation of benefit(s) (expected attainment of the specified objectives), and costs (the resource cost(s) of the selected options). By varying the selection of those options, it is possible to consider alternate points within the space of possible development decisions, each with their own costs and benefits.

Ability to trace the purpose of the selected risk- or defect-reducing options. Within the DDP model these trace to the risks or defects they reduce, which in turn trace to the threatened objectives. As a result, a development activity (e.g., a preventative measure such as training) can be traced to the defects whose prevalence it decreases, and in turn to the increase in expected attainment of objectives that will accrue. The DDP model is capable of showing the net benefit of early-lifecycle risk-reduction measures in terms of both their improvement to the ultimate quality of the final product, and their net reduction in development costs (when the prevention or early detection of problems saves much more costly repair should they be discovered later in the lifecycle). (For examples drawn from the software assurance domain, see [5]).

In order to attain the above benefits, it is necessary to (i) gather the information (e.g., the information on the risk-reducing effectiveness of the various options), (ii) perform the appropriate calculations with that information, and (iii) present the results in such a way as to support decision-making over those options. We discuss these next.

2.2 Information elicitation

The DDP approach requires gathering information that would not normally be asked for, in particular, the information on the risk-reducing effectiveness of the various options. In a traditional risk assessment, what would be asked for would be the risk status of the artifact itself (e.g., the reliability anticipated of an appropriately qualified part purchased from a trusted vendor). Instead, the DDP approach calls for asking for the risk-reducing effectiveness of the steps that went into that artifact's construction and qualification. *It takes additional time and effort to gather this information.* In practice we perform detailed information gathering on just those aspects pertinent to the case at hand – typically those which represent the novel aspects of the design being studied.

For example, one of our studies was applied to determine the risks of using a terrestrial technology in a space setting, with the objective being to determine the set of qualification tests needed to perform on the candidate technology. The focus of concern was on whether the technology would survive the temperature swings if used on the surface of Mars

(where the day/night temperature swings are quite large, and, relative to Earth, have unusual low points). Thus we spend the majority of the time and effort identifying potential failure modes that would be caused by temperature swings, and assessing the effectiveness of various prevention and detection measures against those failure modes. We spent relatively less time and effort on other failure modes (e.g., that the technology would exceed the mass and size limitations) which, while no less important, were felt to be better understood. To get an idea of scale, this study (a fairly typical one) took into account 50 distinct objectives, 31 defects, and 58 options for development practices. Approximately 500 quantitative impact links connected defects to objectives, and 300 quantitative effect links connected options to defects.

Gathering this amount of information is typically done in several half-day sessions, during which experts representing each of the disciplines involved are simultaneously present. The DDP software helps by allowing for on-the-fly capture of information as expressed by the experts. For details, see [6].

2.2 Risk calculations

Calculation over the accumulated information is done by the DDP software. Because its risk model is relatively simple (albeit voluminous – with hundreds or thousands of items), for a given selection of options it is fast to compute the cost and benefit – typically less than one second on a laptop computer.

2.3 Decision making

The purpose of gathering and calculating with this information is to help in decision making. The primary area of decisions that this approach supports are those of *which options to select*, the issue being that in most circumstances the cost were *all* the risk-reducing options to be selected would far exceed the available resources, hence the need for judicious selection among those options. On occasion it becomes apparent that, for a limited amount of resources and a set of high expectations for objective attainment (both of which are recurrent phenomena in our setting), there does not exist *any* selection from among the options that stays within the resource limits and achieves the requisite levels of objective attainment. In such situations another option is to *discard some of the objectives* (i.e., reduce expectations). In our studies of technologies, this may correspond to limiting somewhat the intended range of application for the technology in question. At the project/mission level, this may correspond to discarding, or downgrading, some of the overall objectives.

The key to enabling this kind of decision-making is the calculation of cost and benefit of a given selection of options. As mentioned above, this is speedily performed by the DDP software. The DDP software offers several ways of visualizing the information calculated from a DDP model, discussed next.

2.3 Information visualization for decision-makers

Several forms of visualizing the results are supported to enable human decision makers understand the status of a

given selection of options: in addition to displaying the overall figures of cost and benefit, DDP also can display the status of individual elements – for each objective, the degree to which its attainment is detracted from by the extant risks; for each risk, the sum total reduction in objective attainment attributable to it; for each option, the increase in objective attainment that accrues from the selection of that option (because it decreases risks' severities and/or likelihoods, and hence leads to increased objective attainment).

The DDP software employs straightforward bar-chart presentations of these sets of information for a single selection of risk mitigation options. They are also effective for comparison of a *pair* of risk mitigation option selections. Figure 2 is an example in which the status of 31 risks is shown as a series of bars: the height of the green bars indicate the initial risk levels (were no mitigation options selected), the height of the red bars indicate risk levels for one of the options, the yellow shows where that option's risk is lower than the other option, the black shows where it is higher than the other option.




Figure 2 - bar chart risk comparison

For detailed scrutiny of several alternate risk mitigation option selections we switch to use of Kiviat charts (a.k.a. "Spider" charts). Figure 3 shows the individual risks' status corresponding to three selections of mitigation options (the points on the radial lines joined by the blue, purple and black line segments respectively), plus the completely unmitigated risk status (indicated by the green line segments). Further out along a radial line denotes a larger value (in the DDP tool, a numerical scale is located on a separate portion of the display).




Figure 3 - risk status display of several option selections

Hand-selection and detailed scrutiny of mitigation options is not necessarily the most effective way of arriving at an

optimal selection of mitigations. The challenges stem from the sheer number of possible selections (for 58 mitigation options, there are 2^{58} , approximately 10^{17} , possible selections), and the intertwined nature among the objectives, risks and mitigation options (a recurring phenomena). In response, we have incorporated heuristic search into the DDP software, using which we can automatically locate near-optimal selections of mitigation options. For example, for a given cost bound, we use heuristic search to locate the selection of mitigations that maximizes attainment of objectives while costing no more than that cost bound. Our current implementation uses simulated annealing [7] as the heuristic search mechanism. We have also experimented with Genetic Algorithms, and a form of machine learning. Generally, these take several minutes to arrive at a reasonably near-optimal solution. For more details of our studies of search techniques in this setting, see [8].

In order to gain an understanding of the overall cost/benefit tradespace, we use a *series* of heuristic searches spaced across the cost spectrum. This is computationally expensive, taking several hours on a typical DDP dataset. The plot, for the same dataset as used in the previous figures, is seen in Figure 4. Each of the approximately 300,000 individual points in the black “cloud” (a single point color is used throughout) corresponds to a selection of mitigation options. For any given cost (position along the horizontal axis), the optimal solutions (selection of mitigation options) are those that lie closest to the top of the chart. Hence the upper boundary of the black cloud indicates the optimal frontier - also referred to as the “Pareto front” [9]. This is very revealing as to what levels of objective attainment can be had for varying levels of funding. Apparent are phenomena such as a “law of diminishing returns”, where the frontier levels off – increasing the funding only marginally increases the possible benefit. Conversely, at low cost levels, only modest increases in funding lead to significant improvement in the

benefit that can be attained.

A wealth of information underpins these charts. For example, in Figure 4 each of the hundreds of thousands of points represents a distinct selection from among the 58 mitigation options. We have explored mixtures of further visualization and computation to gain additional insights from this data. In [10] we describe using metrics of “difference” between selections (based on the differences between the options they each employ) to identify interestingly dissimilar designs and clusters of designs. In [11] we describe a visualization that highlights the contributions of individual mitigation options within the cost/benefit tradespace. To date we have preferred to build these capabilities into the DDP tool itself, but have also begun to experiment with utilizing sophisticated information visualization capabilities that others have built, e.g., we have experimented with ATSV (the ARL Trade Space Visualizer) tool from the University of Pennsylvania [12] to examine large sets of selections.

3. EVOLUTIONS OF THE DDP MODEL TO EXTEND ITS REPRESENTATION OF DESIGN

Our core model, discussed above, has a relatively simple representation of risk (equivalently, defects) – atomic entities that were independently related to the objectives they detracted from, and the mitigation options (e.g., standards, tests, analyses) that would reduce those risks. We recognized the need to expand our model to better represent designs. We have explored two paths towards this end: elaborations within the DDP model itself, and interfaces to complementary models.

3.1 Elaborations within the DDP model - realization




Figure 4 - cost/benefit tradespace

The key step we have taken in this direction is introduction of more structure into the representation of the risks in our model. We adopted the notions of fault trees [13], in particular the “And” and “Or” gates from which fault trees are constructed. In place of the simple layer of atomic risks, we now have fault tree structures, as sketched abstractly in

Figure 5.
The root

nodes of these fault trees represent failures with potential for impact on objectives.




Figure 5 - fault trees within the DDP model

Their likelihoods are computed from the structure of the tree, and the likelihoods of the leaf nodes. Thus in DDP, the links that connect risks to objectives now connect the roots of fault trees to the objectives, where occurrence of those root events would detract from the attainment of those objectives. (On occasion it is also appropriate to link non-root nodes of a fault tree to objectives). The unique feature of DDP – its explicit treatment of mitigation options linked to the risks they effect (generally, reduce) – is retained in this elaboration; the mitigation options are linked to the appropriate nodes in a fault tree, depending on their nature. When a mitigation is of the kind that reduces the *severity* of a risk, it is connected to the root of the fault tree, and serves to diminish the severity of the impacts that fault has on objectives. When a mitigation is of the kind that reduces by prevention the likelihood of adverse events (faults), it is linked to the leaf-nodes representing those faults, and (when selected) serves to decrease their likelihoods of occurrence. Since the likelihoods of the root nodes are calculated from the structure of the fault tree, and the likelihoods of the leaf nodes of the tree, effecting the likelihoods of leaf nodes is a way of effecting the likelihoods of the root nodes, and therefore their overall expected detraction from objectives. Finally, when a mitigation is of the kind that reduces by detection and repair the likelihood of adverse events (faults), it is linked to whatever level in the fault tree that detection is applied. For example, if the leaves of a fault tree represent potential faults in individual components, then unit testing of a component (followed by repair of any problems found) will yield a net decrease in the likelihood of such faults remaining in that component. Similarly, if a system test is performed, any problems it finds will be repaired in the components themselves, so the net result is again a net decrease in the likelihood of the faults remaining in those components, and hence in the system. For further details, see [14].

The considerations that stemmed from a multitude of risk mitigations options in the original kind of DDP models

continue to apply here – namely, the challenge of determining cost-effective selections. Fortunately, having extended the DDP calculations of risk to perform the fault tree likelihood calculations, DDP’s other mechanisms continue to apply. In particular, optimization utilizing heuristic search is still available. The introduction of fault trees into the risk layer complicates the internal DDP calculations (of risk likelihoods, and therefore of benefit, measured as the sum of objective attainment levels), but from the point of view of the optimizer, a DDP model continues to be utilized to compute the cost and benefit of a selection of mitigation options, just as before.

We are experimenting with ways to give the users detailed insight into the individual contributions of the elements of the fault tree structures themselves. One of these uses color-coding of the fault tree elements to indicate likelihoods. Figure 6 illustrates this with a “zoom-in” to a small portion of a larger fault tree. The colors are allocated on a red-to-green spectrum, with red representing the higher likelihood values, green the lower ones (in the actual DDP tool, a numerical scale is located on a separate portion of the display). Thus we can see that the red-colored “Or” gate is one of the larger contributors to likelihood in the visible portion of this structure; in turn, its leftmost child is an orange color, whereas the other two children are colored with greenish tints – hence we can see that its leftmost child is the most likely of the three.




Figure 6 - zoom in on fault tree elements

3.1 Elaborations within the DDP model - utilization

Having added fault trees to the DDP model, we use them to expand DDP’s ability to represent design and development alternatives.

The most obvious of these is to represent the structure of a proposed design in the manner for which fault trees are traditionally applied. That is, the manner by which fault(s) in combination give rise to failures (e.g., as in systems with built-in redundancy to make them fault tolerant).

We also have found a use for them to represent design *alternatives*. To represent a design alternative, we add both the purpose that design fulfils as a risk (which more intuitively could be thought of as a problem), and the design option itself as a mitigation option that, if selected, mitigates (solves) that risk.

For example, if we have need for electrical power, then “lack of electrical power” is added as a DDP risk, and the design alternatives for providing power are added as DDP mitigations, each linked to that same risk. Thus if we didn’t select any of those design alternatives, the “lack of electrical power” risk would be unmitigated, and presumably detract from attainment of all the objectives requiring electrical

power. If we select *one* of those design options, the “lack of electrical power” risk will be mitigated. If we select more than one of those design options, the “lack of electrical power” risk will be doubly (maybe with no net improvement over singly) mitigated. If one power source is sufficient, and since, presumably, the *costs* of selecting more than one will accumulate, the usual way of locating cost-effective mitigation option selections will home in on the choices among single selections of power sources – i.e., handle them as alternatives.

We can also represent *risks specific to a given design alternative* (for example, some kinds of electrical power sources may be potential sources of dangerous levels of heat, which would pose risks to other components in the design). To do this, we include those risks as DDP risks, but give them an initial likelihood of occurrence of zero. Then, each design option whose selection would trigger that risk is linked to it by a DDP effect link which *increases*, rather than decreases, the risk likelihood. DDP effect links that increase risks were already present within the DDP model, intended to represent things like tests with potentially detrimental side effects (e.g., vibrating a piece of equipment in a shake test may be an effective way of revealing the presence of certain defects, but may itself create damage, especially if done incorrectly). Here we make use of this same capability to turn “on” risks.

The net result is that we are able to represent within the same DDP framework both development practices (preventions, tests, analyses etc) as options, and design alternatives as options.

To date we have dealt primarily with risk factors of expected loss, as described herein. We have recently incorporated the notion of *ranges* of values (of defect likelihoods, impacts, effects and costs). Using these we can calculate best-case and worst-case extremes. For example, the best-case extreme is calculated assuming the low-ends of defect likelihoods, the low-ends of impact values, the low-ends of costs, and the *high*-ends of effect values (high-ends of these because the higher effect value reduces a defect more, and so is the best-case). However, this incurs the extra time and effort to elicit ranges (rather than a single value). We are able to minimize this by first eliciting just the expected values, performing sensitivity analysis on them, and eliciting ranges for only those values to which the overall computations are most sensitive.

As pointed out by a reviewer of this paper another avenue worthy of exploration is of time-dependent failure rates, as would be needed for calculations of availability and reliability. This represents an area of future work.

3.2 Interfaces to complementary models

We have also explored an approach in which we interface to other models, ones that already have in place structures (such as fault trees and event trees) but which lack an explicit representation of options. This allows us to call upon the power of those other models to represent and reason over those structures, while retaining our ability to explore (compare, optimize over, etc) options and their combinations.

Our most extensive such study to date has been connecting DDP with the dynamic fault tree tool

Galileo/ASSAP, developed at the University of Virginia: [15] and [16]. Fault trees are constructed in DDP, written out in the textual format that the Galileo/ASSAP tool understands, opened up in Galileo/ASSAP for it to evaluate likelihoods, and the evaluation results written back into a file for DDP to read. We have made this information flow an automated process.

The net result is that DDP is able to utilize the fault tree calculations already implemented in Galileo/ASSAP. In the case of large fault trees, with numerous “shared” nodes (where the same fault plays a role in multiple places within a tree), the Galileo/ASSAP implementation is much more efficient than DDP’s, making use of binary decision diagrams (BDDs) to solve static sub-trees [17] rather than the naive DDP implementation. Meanwhile, the DDP contribution to this pairing of the two tools is its representation of the risk reducing mitigations as options, and the built-in simulated annealing optimizer can be brought to bear.

In ongoing work we are following a similar approach to connect DDP to another tool, one that represents event-consequence trees. Again, the purpose is to take advantage of the pre-built analysis capabilities offered by that other tool, in conjunction with DDP’s explicit representation of risk mitigations as options.

7. ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration and funded through the internal Research and Technology Development program, NASA’s Office of Safety and Mission Assurance, and the former NASA Codes R and T. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

We especially thank Chester Borden, Yuri Gawdiak, Kenneth Hicks, Jelly Moran, Irem Tumer & Stephen Prusha for making this work possible, and Profs. Joanne Dugan and Kevin Sullivan of the University of Virginia for their cooperation in the DDP-Galileo/ASSAP connection work.

REFERENCES

1. S.L. Cornford, “Managing Risk as a Resource using the Defect Detection and Prevention process”, *4th International Conference on Probabilistic Safety Assessment and Management*, (September) 1998, New York City, NY, International Association for Probabilistic Safety Assessment and Management.
2. M.S. Feather, S.L. Cornford & M. Gibbel, “Scalable Mechanisms for Requirements Interaction Management”, *Proceedings, 4th IEEE International Conference on Requirements Engineering*, (June) 2000, pp 119-129.
3. T. Ferguson, “*Mathematical Statistics: A Decision Theoretic Approach*”, Academic Press, 1967.

4. "Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners, version 1.1", prepared for the Office of Safety and Mission Assurance, NASA HQ, Washington, DC, August 2002; <http://www.hq.nasa.gov/office/codeq/doctree/praguide.pdf>
5. M.S. Feather, B. Sigal, S.L. Cornford & P. Hutchinson, "Incorporating Cost-Benefit Analyses into Software Assurance Planning", *Proceedings, 26th IEEE/NASA Software Engineering Workshop*, Greenbelt, Maryland November 27-29 2001. IEEE Computer Society, pp 62-68.
6. M.S. Feather, S.L. Cornford, K.A. Hicks & K.A. Johnson, "Applications of tool support for risk-informed requirements reasoning" *Computer Systems Science and Engineering* ([CRL Publishing Ltd](http://www.crlpublishing.com)); 20(1): 5-17, January 2005.
7. S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi, "Optimization by simulated annealing". *Science*, Number 4598, 13 May 1983, 220, 4598:671-680.
8. S.L. Cornford, M.S. Feather, J.R. Dunphy, J. Salcedo & T. Menzies, "Optimizing Spacecraft Design – Optimization Engine Development: Progress and Plans", *Proceedings of the 2003 IEEE Aerospace Conference* (Big Sky, Montana, March 2003), pp 7-3361 – 7-3368.
9. P. Sen & J-B. Yang, *Multiple Criteria Decision Support in Engineering Design*, Springer-Verlag, 1998.
10. M.S. Feather, J. Kiper, & S. Kalafat, "Combining Heuristic Search, Visualization and Data Mining for Exploration of System Design Spaces", *14th Annual International Symposium Proceedings of INCOSE 2004* (Toulouse, France, June 20-24 2004).
11. M.S. Feather, "Towards Cost-Effective Reliability through Visualization of the Reliability Option Space" *Proceedings of the 2004 Annual Reliability and Maintainability Symposium (RAMS)*; Los Angeles, CA, Jan. 2004, pp. 546-552.
12. G.M. Stump, M. Yukish, T.W. Simpson & L. Bennett, "Multidimensional Visualization and Its Application to a Design by Shopping Paradigm", *Proceedings 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*. (Atlanta, GA, 2002), AIAA, AIAA-2002-5622.
13. W.E. Vesely, F.F. Goldberg, N.H. Roberts & D.F. Haasl, "Fault Tree Handbook", U.S. Nuclear Regulatory Commission NUREG-0492, 1981.
14. M.S. Feather, "Towards a Unified Approach to the Representation of, and Reasoning with, Probabilistic Risk Information about Software and its System Interface", *15th IEEE International Symposium on Software Reliability Engineering*, Saint-Malo, Bretagne, France, 2-5 November 2004, pp 391-402.
15. J.B. Dugan, S.J. Bavuso & M. Boyd, "Dynamic Fault Tree Models for Fault Tolerant Computer Systems," *IEEE Trans. on Reliability*, 41(3), Pages 363-377, Sept. 1992.
16. K.J. Sullivan, J.B. Dugan & D. Coppit, "The Galileo Fault Tree Analysis Tool", *Proceedings of the 29th International Conference on Fault-Tolerant Computing (FTCS-29)*, 1999.
17. S.A. Doyle & J.B. Dugan, "Dependability Assessment using Binary Decision Diagrams (BDDs)", *25th Annual International Symposium on Fault-Tolerant Computing*, Pasadena, California, 27-30 June 1995.

BIOGRAPHIES

Martin S. Feather, PhD,
Jet Propulsion Laboratory, California Institute of Technology,
4800 Oak Grove Drive,
Pasadena, CA 91109-8099 USA

e-mail: Martin.S.Feather@jpl.nasa.gov

Martin S. Feather is a Principal in the Software Quality Assurance group at JPL. He works on developing research ideas and maturing them into practice, with particular interests in the areas of software validation (analysis, test automation, V&V techniques) and of early phase requirements engineering and risk management. He obtained his BA and MA degrees in mathematics and computer science from Cambridge University, England, and his PhD degree in artificial intelligence from the University of Edinburgh, Scotland. For further details, see <http://eis.jpl.nasa.gov/~mfeather>

Steven L. Cornford, PhD
Jet Propulsion Laboratory, California Institute of Technology,
4800 Oak Grove Drive,
Pasadena, CA 91109-8099 USA

e-mail: Steven.L.Cornford@jpl.nasa.gov

Steven Cornford is a Senior Engineer in the Strategic Systems Technology Program Office at NASA's Jet Propulsion Laboratory. He graduated from UC Berkeley with undergraduate degrees in Mathematics and Physics and received his doctorate in Physics from Texas A&M University in 1992. Since coming to JPL he focused his early efforts at JPL on establishing a quantitative basis for environmental test program selection and implementation. As Payload Reliability Assurance Program Element Manager, this evolved into establishing a quantitative basis for evaluating the effectiveness of overall reliability and test programs as well as performing residual risk assessments of new technologies. This has resulted in the Defect Detection and Prevention (DDP) process is the motivation for this paper. He received the NASA Exceptional Service Medal in 1997 for his efforts to date. He has been an instrument system engineer, a test-bed Cognizant Engineer and is currently involved with improving JPL's technology infusion processes as well as the Principal Investigator for the development and implementation of the DDP software tool.